

Robust Structured Low-Rank Approximation on the Grassmannian

Clemens Hage and Martin Kleinsteuber

Department of Electrical Engineering and Information Technology
Technische Universität München
Arcisstr. 21, 80333 Munich, Germany
{hage,kleinsteuber}@tum.de,
<http://www.gol.ei.tum.de>

Abstract. Over the past years Robust PCA has been established as a standard tool for reliable low-rank approximation of matrices in the presence of outliers. Recently, the Robust PCA approach via nuclear norm minimization has been extended to matrices with linear structures which appear in applications such as system identification and data series analysis. At the same time it has been shown how to control the rank of a structured approximation via matrix factorization approaches. The drawbacks of these methods either lie in the lack of robustness against outliers or in their static nature of repeated batch-processing. We present a Robust Structured Low-Rank Approximation method on the Grassmannian that on the one hand allows for fast re-initialization in an online setting due to subspace identification with manifolds, and that is robust against outliers due to a smooth approximation of the ℓ_p -norm cost function on the other hand. The method is evaluated in online time series forecasting tasks on simulated and real-world data.

1 Introduction

Many applications such as system identification and time series analysis motivate the problem of Structured Low-Rank Approximation (SLRA). While common low-rank approximations like PCA aim to find a low-dimensional subspace to represent high-dimensional data optimally with respect to some norm or divergence, in the structured case this problem is extended by the additional constraint that the low-rank approximation has to meet a certain linear structure (Hankel, Toeplitz, Sylvester).

For the prominent case of Hankel matrices a method dubbed Singular Spectrum Analysis (SSA) [4] has been presented, which performs the simplest way of SLRA in that it computes a low-rank approximation of a (Hankel-) structured matrix followed by a so-called diagonal averaging step, which is the projection onto the space of Hankel matrices. An obvious drawback of this method is that this projection can destroy the low-rank property established before. The method of Cadzow [6] alternates between these two steps until the algorithm converges to a solution that is indeed low-rank and structured. However, as Chu et al. [8] and Markovsky [12] state, this solution can be far away from the initialization with no guarantees of finding an actually meaningful approximation to the data. Recently, Ishteva et al. [11] have proposed a factorization

approach with a cost function that joints the structural and low-rank constraint. The dimension of the two matrix factors is an upper bound on the rank of the approximation and its structure is enforced with a side condition. The approximation is fitted to the data according to an ℓ_2 -norm, although it is well known that low-rank approximations of this kind can be vulnerable against outliers. For the unstructured case this has been the major incentive to move from PCA to robustified PCA methods such as the Robust PCA method by Candes et al. [7] which recovers a subspace in the presence of sparse outliers of great magnitude. Ayazoglu et al. [1] have proposed to extend this concept to structured matrices by introducing additional Lagrangian multipliers. Their method is called Structured Robust PCA (SRPCA), and its performance is demonstrated in visual applications like Target Location Prediction, Tracklet Matching (both matrix completion problems) and Outlier Removal from trajectories, which can be interpreted as outlier identification through robust subspace estimation.

One of the drawbacks of many Robust PCA approaches and thus also the SRPCA approach is their batch-processing nature. In an online setting the algorithm needs to be re-run from scratch if the data set grows or changes over time. This can be alleviated by factorizing the low-rank approximation of data into an orthogonal matrix representing the subspace and a coefficient matrix containing the coordinates of the currently observed data in this subspace, cf. [2,9] for the unstructured case. Whenever new data comes in, it is possible to initialize the subspace optimization with the previous estimate and to update both the subspace and the coordinates to the new data. Obviously, whenever the subspace does not change significantly this saves computational effort compared to a random initialization. We will firstly derive a batch algorithm for Robust Structured Low-Rank Approximation on the Grassmannian and then outline how to process structured data online in an efficient way. We illustrate the performance of the proposed algorithm on several time series forecasting tasks.

2 Robust Structured Low-Rank Approximation on the Grassmannian using a smoothed ℓ_p -norm cost function

2.1 Low-rank and sparsity constraints in the unstructured case

Low-rank approximation of data is a well-studied problem, cf. [13] for a recent overview. In the past years a trend can be seen towards Robust PCA methods that are tolerant against outliers in the data. An often considered data model is $\mathbf{X} = \mathbf{L} + \mathbf{S}$, where the input data $\mathbf{X} \in \mathbb{R}^{m \times n}$ is assumed to be composed of a low-rank part \mathbf{L} with $\text{rank}(\mathbf{L}) \leq k$ and a sparse matrix \mathbf{S} that contains few non-zero entries. While in [7] the low-rank constraint is enforced via minimization of the nuclear norm, in this work we will consider the Robust PCA setting on the *Grassmannian*

$$\text{Gr}_{k,m} := \{\mathbf{P} \in \mathbb{R}^{m \times m} | \mathbf{P} = \mathbf{U}\mathbf{U}^\top, \mathbf{U} \in \text{St}_{k,m}\}. \quad (1)$$

The subspace is hereby represented by an element \mathbf{U} of the *Stiefel* manifold $\text{St}_{k,m} = \{\mathbf{U} \in \mathbb{R}^{m \times k} | \mathbf{U}^\top \mathbf{U} = \mathbf{I}_k\}$, where \mathbf{I}_k is the $(k \times k)$ -identity matrix. The low-rank approximation is then $\mathbf{L} = \mathbf{U}\mathbf{Y}$ with $\mathbf{Y} \in \mathbb{R}^{k \times n}$.

In contrast to the ℓ_2 -norm in common PCA, robust approaches use relaxations of the ℓ_0 -norm for fitting the low-rank approximation to the data, as gross outliers might otherwise distort the estimation. The smoothed ℓ_p -norm

$$h_\mu: \mathbb{R}^{m \times n} \rightarrow \mathbb{R}^+, \quad X \mapsto \sum_{j=1}^n \sum_{i=1}^m \left(x_{ij}^2 + \mu \right)^{\frac{p}{2}}, \quad 0 < p < 1 \quad (2)$$

presented in [9] behaves similarly to the ℓ_0 -norm for sparse outliers, but treats small additive Gaussian noise like an ℓ_2 norm. A Robust Low-Rank Approximation problem using the sparsifying function (2) writes as

$$(\hat{\mathbf{U}}, \hat{\mathbf{Y}}) = \arg \min_{\mathbf{U} \in \text{St}_{m,k}, \mathbf{Y} \in \mathbb{R}^{k \times n}} h_\mu(\mathbf{X} - \mathbf{U}\mathbf{Y}), \quad \hat{\mathbf{L}} = \hat{\mathbf{U}}\hat{\mathbf{Y}}. \quad (3)$$

The sparse component can be recovered via $\hat{\mathbf{S}} = \mathbf{X} - \hat{\mathbf{L}}$, possibly followed by a thresholding operation to remove residual noise.

2.2 Extension to linear matrix structures and algorithmic description

Besides the low-rank and sparsity decomposition no other assumptions are made on the data in (3). In many applications however, structured matrices like Hankel, Sylvester or Toeplitz matrices play an important role. Following the notation of [11] we denote by \mathcal{S} a linear matrix structure, by $\mathcal{S}(\mathbf{d})$ a structured matrix obtained from a data series \mathbf{d} and by $\mathcal{P}_{\mathcal{S}}(\mathbf{X})$ the orthogonal projection of any (possibly unstructured) matrix \mathbf{X} onto the image of \mathcal{S} w.r.t. the standard inner product. For example, if \mathcal{H} is a Hankel structure the orthogonal projection $\mathcal{P}_{\mathcal{H}}$ is equivalent to the diagonal averaging step in SSA [4].

In order to include the structural constraint we extend the cost function (3) by the side condition $\mathbf{L} \in \mathcal{S}$. This motivates the Lagrangian Multiplier scheme

$$\min_{\mathbf{U} \in \text{St}_{m,k}, \mathbf{Y} \in \mathbb{R}^{k \times n}, \Lambda \in \mathbb{R}^{m \times n}} h_\mu(\mathbf{X} - \mathbf{U}\mathbf{Y}) + \langle \Lambda, \mathbf{U}\mathbf{Y} - \mathcal{P}_{\mathcal{S}}(\mathbf{U}\mathbf{Y}) \rangle + \frac{\rho}{2} \|\mathbf{U}\mathbf{Y} - \mathcal{P}_{\mathcal{S}}(\mathbf{U}\mathbf{Y})\|_F^2. \quad (4)$$

Algorithm 1 outlines the extension of the Robust PCA method from Hage and Kleinstenuber [9] to the case of structured matrices. The algorithm considers a partial observation $\hat{\mathbf{X}}$ of the data with \mathcal{A} defining which entries are actually observed. In each iteration, three steps are performed. Firstly, the subspace estimate is updated. In this realization the subspace is uniquely identified with a Grassmannian projector $\mathbf{P} = \mathbf{U}\mathbf{U}^\top$ with $\mathbf{U} \in \text{St}_{k,m}$. The optimization problem is solved via Conjugate Gradient (CG) descent with backtracking line-search using a QR-decomposition based retraction on the Grassmannian. Once \mathbf{P} and thereby \mathbf{U} have been found the coordinates \mathbf{Y} are updated with a CG method in Euclidean space, such that a new optimum low-rank estimate $\mathbf{L} = \mathbf{U}\mathbf{Y}$ is found. In a third step the Lagrangian multiplier Λ is updated, then ρ is increased and μ is decreased. While μ controls the behavior of the sparsifying function $h_\mu(\cdot)$, the parameter ρ weighs between the data fitting term and the structural side condition. More precisely, as long as ρ is small the Robust PCA term is the leading power and a low-rank approximation is fitted to the data. With increasing ρ the structural condition is more and more enforced until it is the dominating term in the cost function.

In a practical application the optimization can also be terminated if the residual Hankel penalty ϵ , i.e. the Frobenius distance to the next Hankel-structured matrix normalized by the number of matrix entries falls below a certain threshold τ .

3 Efficient Online Time Series Forecasting via Robust SLRA

We have outlined how to use manifold optimization for Robust SLRA on the Grassmannian. In the important case of a Hankel structure SLRA corresponds to identifying an LTI system, cf. [13]. In practical applications, however, an observed system might be time-variant. Or the observed data is not related to a physical system at all but still exhibits repetitive or periodic behavior. The field of Time Series Analysis [3] deals with these signals and numerous autoregressive methods for filtering and forecasting data series exist. In the SLRA context a low-rank Hankel matrix (and thus an LTI) is fitted to the observed data and the future development is extrapolated from this approximation. Thereby, the rank bounds the complexity of the approximation. If the behavior of the data changes over time a new model needs to be determined for each observation instance. For our proposed Grassmannian Robust SLRA method this means that both a new subspace and new coordinates need to be computed. However, when the signal characteristics vary moderately over time it is likely that the new subspace lies close to the previously found one. Therefore, the subspace should not be randomly initialized but rather updated with the new data point, in a similar way as Robust Subspace Tracking ([10], [14]) in the unstructured case. As discussed earlier, however, we do not optimize directly on the space of structured low-rank matrices. Instead we relax the structural constraint, update the subspace and then tighten the structural side condition again by varying the parameter ρ in the cost function.

Algorithm 1 Alternating minimization scheme for Grassmannian Robust Structured Low-Rank Approximation

Initialize:

Choose $\mathbf{X}_0 \in \mathbb{R}^{m \times n}$, s.t. $\mathcal{A}(\mathbf{X}_0) = \hat{\mathbf{X}}$.

Initialize $\mathbf{U}^{(0)}$ randomly or from k left singular vectors of \mathbf{X}_0 .

$$\mathbf{Y}^{(0)} = \mathbf{U}^{(0)\top} \mathbf{X}_0, \mathbf{L}^{(0)} = \mathbf{U}^{(0)} \mathbf{Y}^{(0)}$$

$$\mathbf{P}^{(0)} = \mathbf{U}^{(0)} \mathbf{U}^{(0)\top}$$

Choose $\mu^{(0)}$ and $\mu^{(l)}$, compute $c_\mu = \left(\frac{\mu^{(l)}}{\mu^{(0)}}\right)^{1/(l-1)}$

Choose $\rho^{(0)}$ and $\rho^{(l)}$, compute $c_\rho = \left(\frac{\rho^{(l)}}{\rho^{(0)}}\right)^{1/(l-1)}$

for $i = 1 : l$ **do**

$$\begin{aligned} \mathbf{P}^{(i+1)} = \arg \min_{\mathbf{P} \in \text{Gr}_{k,m}} & h_{\mu^{(i)}}(\hat{\mathbf{X}} - \mathcal{A}(\mathbf{P}\mathbf{L}^{(i)})) \\ & - \langle \Lambda^{(i)}, \mathbf{P}\mathbf{L}^{(i)} - \mathcal{P}_S(\mathbf{P}\mathbf{L}^{(i)}) \rangle \\ & + \frac{\rho^{(i)}}{2} \|\mathbf{P}\mathbf{L}^{(i)} - \mathcal{P}_S(\mathbf{P}\mathbf{L}^{(i)})\|_F^2 \end{aligned}$$

$$\text{find } \mathbf{U}^{(i+1)} \quad \text{s.t.} \quad \mathbf{U}^{(i+1)} \mathbf{U}^{(i+1)\top} = \mathbf{P}^{(i+1)}$$

Subspace Step

$$\begin{aligned} \mathbf{Y}^{(i+1)} = \arg \min_{\mathbf{Y} \in \mathbb{R}^{k \times n}} & h_{\mu^{(i)}}(\hat{\mathbf{X}} - \mathcal{A}(\mathbf{U}^{(i+1)}\mathbf{Y})) \\ & - \langle \Lambda^{(i)}, \mathbf{U}^{(i+1)}\mathbf{Y} - \mathcal{P}_S(\mathbf{U}^{(i+1)}\mathbf{Y}) \rangle \\ & + \frac{\rho^{(i)}}{2} \|\mathbf{U}^{(i+1)}\mathbf{Y} - \mathcal{P}_S(\mathbf{U}^{(i+1)}\mathbf{Y})\|_F^2 \end{aligned}$$

$$\mathbf{L}^{(i+1)} = \mathbf{U}^{(i+1)} \mathbf{Y}^{(i+1)}$$

Coordinate Step

$$\Lambda^{(i+1)} = \Lambda^{(i)} - \rho^{(i)} (\mathbf{L}^{(i+1)} - \mathcal{P}_S(\mathbf{L}^{(i+1)}))$$

Multiplier Update

$$\mu^{(i+1)} = c_\mu \mu^{(i)}, \quad \rho^{(i+1)} = c_\rho \rho^{(i)}$$

$$\epsilon = \frac{1}{mn} \|\mathbf{L}^{(i+1)} - \mathcal{P}_S(\mathbf{L}^{(i+1)})\|_F$$

end for

$$\hat{\mathbf{L}} = \mathcal{P}_S(\mathbf{L}^{(l)}), \quad \hat{\mathbf{S}} = \mathbf{X} - \hat{\mathbf{L}}$$

In Algorithm 2 we describe an Online Time Series Forecasting method based on Robust Structured Low-Rank Approximation on the Grassmannian. The algorithm receives as inputs the time series of data \mathbf{d} to be analyzed as well as the desired order of the system and the forecasting range, i.e. the number of samples to be predicted. Since a Hankel matrix of size $m \times m$ contains $(2m - 1)$ samples of data, the prediction starts at $\mathbf{d}(2m)$. A data vector $\mathbf{x}_{(j)}$ of length $(2m - 1)$ is extracted from the data up to the present position, padded with zeros according to the forecasting range and structured to form a Hankel matrix. The first subspace estimate $\mathbf{U}^{(0)}$ is not initialized randomly but with $\mathbf{U}_{(j-1)}$, the final subspace estimate of the previous set of data samples. Note that the subscript (j) counts the position of the current set of data samples in the data stream while in Algorithm 1 the superscript (i) denoted the iteration count of the alternating minimization steps in the batch process. Accordingly, for each set of data samples at position j the number of alternating minimization steps $I_{(j)}$ for the current estimation must be set beforehand. This number varies between a predefined I_{min} and I_{max} , and the choice is based on the residual Hankel penalty $\epsilon_{(j-1)}$ of the previous iteration. This corresponds to the observation that significant changes in the system behavior lead to higher values for ϵ and require more iterations in the optimization process, whereas less update steps are required if the subspace changes slowly or does not change at all. We start with the previous iterate on the Stiefel manifold and execute the three steps from Algorithm 1 in an alternating manner until convergence. Notice that the cost function parameter μ is fixed here and only the Lagrangian parameter ρ is changed in each iteration to steer between data fit and structure. The forecasting is realized as a robust matrix completion problem, i.e. the respective entries in the lower right corner of \mathbf{X} are considered unobserved entries. Once a structured low-rank estimate has been found, the predicted entries of \mathbf{d} can easily be read from the last entries of $\mathbf{l}_{(j)}$.

Algorithm 2 Online Time Series Forecasting via Grassmannian Robust SLRA

Input: data series $\mathbf{d} \in \mathbb{R}^N$, system order k , analysis dimension m , forecasting range r
for $j = 2m : N$ **do**

 Define $\mathbf{x}_{(j)} = [\mathbf{d}(j - (2m - 1) : j)^\top \mid \mathbf{0}_r^\top]^\top$
 and \mathcal{A} according to forecasting range

 Obtain $\mathbf{X}_{(j)} = \mathcal{H}(\mathbf{x}_{(j)})$

 Initialize $\mathbf{U}^{(0)} = \mathbf{U}_{(j-1)}$

$\mathbf{Y}^{(0)} = \mathbf{U}^{(0)\top} \mathbf{X}_{(j)}$, $\mathbf{L}^{(0)} = \mathbf{U}^{(0)} \mathbf{Y}^{(0)}$

$\mathbf{P}^{(0)} = \mathbf{U}^{(0)} \mathbf{U}^{(0)\top}$

 Select number of iterations $I_{(j)}$

 Choose $\rho^{(0)}$ and $\rho^{(l)}$

 Compute $c_{\rho(j)} = \left(\frac{\rho^{(l)}}{\rho^{(0)}}\right)^{1/(I_{(j)}-1)}$

for $i = 1 : I_{(j)}$ **do**

Subspace Step from Alg. 1

Coordinate Step from Alg. 1

Multiplier Update from Alg. 1

$\rho^{(i+1)} = c_{\rho(j)} \rho^{(i)}$

end for

$\mathbf{l}_{(j)} = \mathcal{P}_{\mathcal{S}}(\mathbf{L}^{(I_{(j)})})$

$\hat{\mathbf{d}}_{(j)}(j + 1 : j + r) = \mathbf{l}_{(j)}(2m : 2m + r - 1)$

end for

4 Experimental results

In a first experiment we evaluate our method on an impulse response prediction task for a noisy observation of a simulated SISO Linear Time Varying (LTV) system. The data




Fig. 1. Forecasting of 3 samples of a SISO-LTV impulse response with additive noise and outliers

is generated via

$$\begin{aligned}\dot{\mathbf{x}}(t+1) &= \mathbf{A}(t)\mathbf{x}(t) + \mathbf{b}^\top \mathbf{u}(t), \quad \mathbf{A}(t) = e^{0.001t\mathbf{Z}}, \quad \mathbf{Z}^\top = -\mathbf{Z} \\ \mathbf{y}(t) &= \mathbf{c}^\top \mathbf{x}(t) + \mathbf{n}(t)\end{aligned}$$

with \mathbf{b}, \mathbf{c} being uniform random vectors and \mathbf{Z} being a random skew-symmetric matrix. The additive observation noise $\mathbf{n}(t)$ contains of two parts, Gaussian noise with $\sigma = 0.01$ and randomly appearing (rate 0.05) salt and pepper noise samples that take on values of ± 0.5 . The degree of the system is chosen as $k = 5$ and we generate the impulse response of the system for 300 samples. Our method is compared to the SLRA method from [11] and both algorithms are implemented in *MATLAB* on a desktop computer. We predict three time steps into the future from an observation of $2m - 1$ samples, and the parameters are empirically chosen as $m = 20$, $\rho \in [10^{-6}, 10]$ (both methods), $p = 0.5$, $\mu = 0.005$, $\tau = 5 \times 10^{-4}$. The SLRA method is randomly initialized in each step, converges within 30 iterations and requires about 0.4 seconds. The iteration number of our method varies between $I_{\min} = 16$ and $I_{\max} = 128$ iterations with an average of 22 iterations that add up to 0.7 seconds per forecasting step. The forecasting results in Figure 1 indicate that both methods are able to cope with the Gaussian noise quite well and predict the system behavior quite reliably, but the spurious outliers introduce errors in the SLRA extrapolation due to the ℓ_2 -error weighting. Our proposed method is much more robust at the price of a higher computational effort. However, due to the beneficial subspace initialization the computation time is still competitive.

In a second experiment we compare our method on real-world data with SLRA and the *forecast* routine in *MATLAB* with a 12-month-seasonal ARIMA(0,1,1) model¹. The time series is the well-known *Airline Passenger* dataset from [3] normalized to the range $[0, 1]$. The upper bound on the rank of the approximation is chosen as $k = 8$, and we forecast 6 samples from $2m - 1$ samples with $m = 18$, which corresponds to projecting the monthly amount of passengers half a year into the future from observing the past three years. Figure 2 shows that all three methods succeed in forecasting the data, with average absolute deviations of 0.060 for SLRA, 0.036 for ARIMA and 0.044 for our method. On average, the ARIMA implementation requires 1.3s, SLRA 0.7s and

¹ <http://mathworks.com/help/econ/forecast-airline-passenger-counts.html>




Fig. 2. Six month forecast of monthly airline passenger data from the years 1952-1960 based on 3 year observation period.




Fig. 3. Six month forecast of monthly airline passenger data from the years 1996-2009 based on 3 year observation period.

our method ($I_{min} = 8, I_{max} = 64$, 12 iterations on average) is the fastest with 0.3s. The popularity of this well-known but also well-behaving dataset has inspired us to perform another experiment on airline passenger data. We have obtained the system-wide (domestic and international) number of passenger emplanements in the USA for the years 1996 – 2014 from the American Bureau of Transportation Statistics [5]. Due to the dramatic developments in the year 2001 this data is obviously more challenging. Figure 3 shows the dataset and the six month forecasts of the three compared methods with the experimental setup as before. The seasonal ARIMA model copes best with the challenging conditions (average absolute error of 0.086), but it needs to be noticed that the actual seasonality is known a priori while SLRA and our method do not have this information. As before, the SLRA method is able to forecast data reliably under good conditions but suffers from the gross outliers, resulting in an average absolute error of 0.182. Finally, our method shows more robustness against gross outliers (average absolute error of 0.102), although in this real-world example the low-rank and sparse data model is not exactly met.

5 Conclusion

We have presented a novel method for Robust Structured Low-Rank Approximation on the Grassmannian. Using an approximated ℓ_p -norm, the method robustly fits an approximation of upper-bounded rank and linear structure to the given data. For the special case of a Hankel structure we have furthermore shown how to use the developed concept for Robust Online Time Series Forecasting. We have shown how to benefit from the manifold setting in online processing, as we can increase the efficiency by re-using the previously identified subspace. Experimental results show that our method performs effectively and efficiently in simulated and real-world applications.

References

1. Ayazoglu, M., Sznaiar, M., Camps, O.I.: Fast algorithms for structured robust principal component analysis. In: Computer Vision and Pattern Recognition (CVPR), 2012 IEEE Conference on. pp. 1704–1711. IEEE (2012)
2. Balzano, L., Nowak, R., Recht, B.: Online identification and tracking of subspaces from highly incomplete information. In: Allerton Conference on Communication, Control, and Computing. pp. 704–711 (2010)
3. Box, G.E.P., Jenkins, G.M.: Time series analysis: forecasting and control, revised ed. Holden-Day (1976)
4. Broomhead, D.S., King, G.P.: Extracting qualitative dynamics from experimental data. *Physica D: Nonlinear Phenomena* 20(2), 217–236 (1986)
5. Bureau of Transportation Statistics, United States Department of Transportation: US Air Carrier travel statistics (2015), <http://www.rita.dot.gov/bts>
6. Cadzow, J.A.: Signal enhancement-a composite property mapping algorithm. *Acoustics, Speech and Signal Processing, IEEE Transactions on* 36(1), 49–62 (1988)
7. Candès, E., Li, X., Ma, Y., Wright, J.: Robust principal component analysis? *Journal of ACM* 58(3), 1–37 (2011)
8. Chu, M.T., Funderlic, R.E., Plemmons, R.J.: Structured low rank approximation. *Linear Algebra and Its Applications* 366, 157–172 (2003)
9. Hage, C., Kleinstüber, M.: Robust PCA and subspace tracking from incomplete observations using ℓ_0 -surrogates. *Computational Statistics* pp. 1–21 (2014)
10. He, J., Balzano, L., Szlam, A.: Incremental gradient on the Grassmannian for online foreground and background separation in subsampled video. In: Computer Vision and Pattern Recognition. pp. 1568–1575 (2012)
11. Ishteva, M., Usevich, K., Markovsky, I.: Factorization approach to structured low-rank approximation with applications. *SIAM Journal on Matrix Analysis and Applications* 35(3), 1180–1204 (2014)
12. Markovsky, I.: Structured low-rank approximation and its applications. *Automatica* 44(4), 891–909 (2008)
13. Markovsky, I.: Low-Rank Approximation - Algorithms, Implementation, Applications. Springer (2014)
14. Seidel, F., Hage, C., Kleinstüber, M.: pROST : A Smoothed ℓ_p -norm Robust Online Subspace Tracking Method for Realtime Background Subtraction in Video. *Machine Vision and Applications* 25(5), 1227–1240 (July 2014)